## Modeling & Simulation

NAS experience in advanced modeling and simulation technologies is essential to creating large-scale simulations supporting critical engineering and design decisions. Our fundamental research advances core numerical methods, algorithms, and physical models for aeronautics, space technologies, heliophysics, and Earth sciences. We also conduct high-fidelity numerical simulation studies that enable NASA engineers to assess and improve the performance, reliability, and safety of aircraft, launch environments, and space exploration vehicles.

For example, in support of NASA's project to design the heavy-lift Space Launch System (SLS), NAS experts used an innovative combination of computational fluid dynamics data and high-fidelity modeling to perform critical simulations and analyze the pressure loads created by the rocket engine plume, helping to inform proposed redesigns of the SLS launch site infrastructure.

Supercomputing for Today and Tomorrow

The NAS approach to real-world supercomputing combines the latest technologies with multi-level security and 24x7 support to give users a reliable and accessible computing experience. This empowers our customers to meet their science and engineering goals while reducing cost, risk, and development time.

Following the sucessful deployment of the modular supercomputing facility that houses our Electra supercomputer, we are planning for an expanded modular enviroment at NASA Ames Research Center with up to five times the capacity of Electra's current configuration. As usage on the Pleiades supercomputer increased more than 20% between 2015 and 2017, these plans for expansion will help keep up with the demands for supercomputing time to support NASA missions.

Together with our industry, university, and government partners, the NAS Division remains committed to supporting current and future agency missions.

Contact Us

To learn more about how our high-end computing resources and customizable services can impact your NASA science and engineering sponsored projects, please contact:

Piyush Mehrotra
NAS Division Chief
(650) 604-5126
piyush.mehrotra@nasa.gov

Bill Thigpen
NAS Systems & Engineering Branch Chief
(650) 604-1061
william.w.thigpen@nasa.gov

NAS High-End Computing Environment Facts

Pleiades

- SGI ICE X cluster connected via InfiniBand in an 11D hypercube topology
- 245,536 cores (Intel Xeon eight-, ten-, twelve-, and fourteen-core processors)
- 7.24 petaflops peak; 5.95 petaflops sustained performance (Nov 2018)
- 935 terabytes total memory
- 64 nodes enhanced with NVIDIA graphics processing units

Electra

- SGI ICE X/HPE E-Cell cluster connected via InfiniBand
- 124,416 cores (Intel Xeon fourteen- and twenty-core processors)
- 8.32 petaflops peak; 5.44 petaflops sustained performance (Nov 2018)
- 589 terabytes total memory

hyperwall

- 128-screen, tiled LCD wall that displays, processes, and shares data
- SGI ICE X system connected to Pleiades via InfiniBand (over 100 gigabytes-per-second connectivity)
- 128 nodes; 2 ten-core (Intel Xeon E5-2680 v2) 280 GHz processors per node
- 128 NVIDIA GeForce GTX 780 Ti graphics processing units
- 8.2 terabytes total system memory; 393 gigabytes GDDR5 graphics memory

Storage

- 40 PB of Lustre "scratch" space
- 1,048-petabyte (1 exabyte) archival storage capacity (with normal 36% compression)
- 3 petabytes of archive disk cache; 98 petabytes of unique data on tape

Networks

- One of the world's largest InfiniBand networks, with over 50 miles of cabling connecting the supercomputers, hyperwall, and storage systems
- Cisco Systems local area network, with 10-gigabit backbone supporting both 1G and 10G hosts
- Hundreds of active ports; several hundred gigabit Ethernet ports; and over 50 10-gigabitsper-second Ethernet ports

For the latest system updates see: http://www.nas.nasa.gov/hecc/resources/environment.html

NASA Advanced Supercomputing Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000
Phone: (650) 604-4502
www.nas.nasa.gov

NASA Advanced Supercomputing Division

NASA Advanced Supercomputing Division

Delivering Powerful and Reliable Computing Solutions for NASA Missions

The NASA Advanced Supercomputing (NAS) Division at Ames Research Center is known worldwide for its innovation and expertise in high-end computing (HEC).

To help achieve exceptional performance and groundbreaking results for NASA's scientific and engineering users, we combine cutting-edge HEC technology and techniques with an emphasis on providing reliable, real-world production capability.

These snapshots of key projects highlight different aspects of our integrated approach to providing custom supercomputing solutions to users across all NASA mission areas.

End-to-End Network Support

NAS network engineers provide custom, end-to-end support services that help users efficiently transfer their massive datasets—often many terabytes in size. The outcome: more computational runs, reduced time-to-solution, and faster turnaround of scientific results.

For example, to support scientists at the Massachusetts Institute of Technology (MIT) who needed to transfer large datasets for the Kepler mission to discover new Earth-like exoplanets, our network experts provided specialized tuning and testing to dramatically increase data transfer rates between MIT and NAS—saving the scientists hundreds of wait-time hours.

Scientific Visualization

The NAS visualization environment connects Pleiades to our in-house hyperwall system via high-speed InfiniBand, forming a unique heterogeneous cluster that enables concurrent visualization—allowing scientists to view their high-resolution simulation results in real time.

With significant, customized support from our "vis" specialists, researchers at NASA Ames were able to explore in 3D the intricate fiber structures of new thermal protection system materials one-tenth the thickness of a human hair. Carbon-fiber felt materials like these are being designed and developed to protect future spacecraft heading beyond low Earth orbit to Mars and other space exploration destinations.

Real-World Supercomputing

Using the Pleiades supercomputer, research scientists supporting NASA's Living with A Star Program produced realistic simulations that show how high-speed jets of plasma develop in Earth's magnetosphere from wind flowing from the Sun's corona. By improving our understanding of how Earth's protective shield interacts with the solar wind, this work may contribute to more accurate forecasts from the impacts of space weather.

As NASA's largest supercomputing resource, Pleiades delivers more than 5.95 petaflops sustained performance. Even though these unique space weather simulations used thousands of processors, Pleiades' large capacity allowed many other mission-critical jobs to run at the same time.

Creative Storage Solutions

Our mass storage capacity, currently 1,040 petabytes (PB) of tape library storage and 2.9 PB of online archive storage, is continually growing to meet user needs. In addition, NAS provides 29 PB of temporary data storage on the HPC systems. Beyond the standard resources provided to users, we provide specialized storage solutions and training to help users manage their massive datasets more efficiently.

For example, researchers utilizing the NASA Earth Exchange (NEX) computing resources used Pleiades to process Landsat global land surface data, which is used by scientists, land use planners, and the public to inform decisions about human health, agriculture, forest management, and more. NAS increased the customized NEX storage infrastructure by 220% to hold more than a petabyte of Landsat data over the next several years.

Code Performance Optimization

NAS code performance experts provide a variety of troubleshooting, porting, and optimization services—ranging from basic to comprehensive—to enhance researchers' productivity and ensure they obtain accurate results in less time.

In one case, NAS experts worked extensively with scientists at NASA's Langley Research Center (LaRC) to improve the performance of code used to help develop supersonic combustion ramjet engines by more than 45 percent. In the process, the LaRC users gained experience in designing their computational fluid dynamics codes to perform well on today's supercomputers.